

Amendments to the Specification:

The numbered paragraphs presented below will replace all prior versions of the paragraphs in the instant application:

[0004] The interaction of the voice prompts and user input is guided by a voice application that in turn is executed by the IVR. Voice applications have been written in script, state code, Java^{[[*]]}™, and voice extensible mark up language (VoiceXML). [[*]]Java™ and all Java™-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc in the United States, other countries or both.

[0008] According to a first aspect of the present invention there is provided an interactive voice response system including: a prompt acquisition component for acquiring an utterance from a user; a speech recognition engine for ~~recognising~~ recognizing a plurality of words from the utterance; a custom server for comparing the actual duration of the utterance with an ideal duration of the ~~recognised~~ recognized words; and a prompt play component for prompting the user as to the speed of delivery of the utterance according to the results of the comparison.

[0009] In this way, data available from a speech recognition engine is used to estimate the speed at which the user is speaking by comparing an ideal duration of the ~~recognised~~ recognized words (as stored with the model of the speech data in the speech recognition engine) with the actual duration of the spoken words.

[0010] Preferably the means for comparing the actual duration of the utterance with an ideal duration of the ~~recognised~~ recognized words include means for acquiring for each word the actual duration of delivery and ideal duration and means for comparing

the differences in actual duration and ideal duration for each word. This solution breaks an utterance down into component words and calculates a difference for each word and then finds the average of all the words. The advantage is that for each ~~recognised~~ word there already exists an ideal duration value in the speech model. The means for acquiring and the means for comparing are defined in the duration custom server.

[0012] More advantageously the means for comparing the actual duration of the utterance with an ideal duration of the ~~recognised~~ words includes calculating an average of the ratio of words as an indication of the speed of delivery of the utterance. Such an average allows a comparison to view the whole picture rather than individual ratios which may on their own distort any conclusion.

[0017] According to a second aspect of the invention there is provided a method in an interactive response system including: acquiring an utterance from a user; ~~recognising~~ a plurality of words from the utterance; comparing the actual duration of the utterance with an ideal duration of the ~~recognised~~ words; and prompting the user as to the speed of delivery of the utterance according to the results of the comparison.

[0018] According to a third aspect of the invention there is provided a computer program product for processing one or more sets of data processing tasks, said computer program product comprising computer program instructions stored on a computer-readable storage medium for, when loaded into a computer and executed, causing a computer to carry out the steps of: acquiring an utterance from a user; ~~recognising~~ a plurality of words from the utterance; comparing the actual duration of the

utterance with an ideal duration of the recognised recognized words; and prompting the user as to the speed of delivery of the utterance according to the results of the comparison.

[0026] According to Figure 1 there is shown a schematic of a voice telephony in which the present invention can be embodied. Voice telephony system 100 comprises an interactive voice response system (IVR) 102 connected to a voice server 116 over a LAN 114. An example of IVR 102 is [[*]]IBM [[*]]WebSphere™ Voice Response 3.1 (WVR) for AIX™ based on IBM [[*]]DirectTalk™ Technology 102. An example of voice server 116 is IBM™ Voice Server. A user uses a telephone 106 to connect with IVR 102 through telephony (PSTN) switch (PABX) 104. IVR 102 uses any one of its three application languages to control a voice interaction. Java application layer 108 uses Java Beans and Java applications to control the IVR 102. State table environment 110 hosts the original DirectTalk application programming language and is based on state table applications and custom servers. VoiceXML application layer 112 uses VoiceXML browsers and VoiceXML applications in Web Servers to control the IVR 100.

[0027] IVR 102 is well-suited for large enterprises or telecommunications businesses. It is scalable, robust and designed for continuous operation 24 hours a day and 7 days a week. IBM WebSphere™ Voice Response 3.1 for AIX™ can support between 12 and 480 concurrent telephone channels on a single system. Multiple systems can be networked together to provide larger configurations. [[*]]AIX™, DirectTalk™, IBM™, pSeries™, and WebSphere™ are trademarks of International Business Machines in the United States, other countries, or both.

[0028] The preferred embodiment uses WebSphere™ Voice Response for AIX™ 3.1 which supports from 1 to 16 E1 or T1 digital trunks on a single IBM pSeries[[*]]™ server with up to 1,500 ports on a single system. Up to 2304 telephony channels using T1 connections or 2880 telephony channels using E1 connections can be supported in a 19" rack. WebSphere™ Voice Response for AIX™ 3.1 requires an IBM AIX™ v 4.3 operating system running on an IBM pSeries™ computer. It supports network connectivity on multiple networks including PSTN, ISDN, CAS, SS7, VoIP networks. The preferred embodiment is concerned with those networks which provide a user identification number with an incoming call e.g. ISDN and SS7.

[0030] The speech recognition engine 118 analyzes input audio using individual pronunciation models for all words in an active vocabulary, including a word representing <silence>. The engine analyzes the audio by fitting it to a mathematical pronunciation model of words in all possible word sequences specified as possible by the vocabulary's language model. The fitting process includes computing a distribution for when each word begins and ends, with the most probable transition points of each distribution reported as the word boundaries. The quality of the mathematical fit between word models and input audio is used together with the language model probability for each word in a particular sequence and several other parameters of the decoding process to compute word scores. During runtime the engine creates these metrics (start time, end time and score) for every word which are passed to the voice custom server along with the recognised recognized word result. An ideal duration time metric of the recognised recognized word result is based on the speech recognition language model. Each phoneme in the language model has an associated ideal duration time and the duration time for a recognised recognized word is the sum of the durations for the phonemes in the recognised recognized word. Normally only the recognised recognized word result is

sent to the IVR but the other metrics are available on demand. In the present embodiment the IVR requests all the above metrics with each utterance it sends to the speech recognition engine.

[0031] Referring to Figure 2, state table "environment 110 of IVR 102 includes: a state table application 202; a duration custom server 204; a voice custom server 206 and an utterance database 208. State table application 202 controls a voice interaction on the IVR 102 when a voice channel from a telephone is opened. The state table application 202 performs application method 300 which is described in relation to Figure 3. Voice custom server 206 provides the interface to the speech recognition engine 118 and the text-to-speech engine 120 on voice server 116. Voice custom server 206 places the results of speech recognition into the utterance database 208 after a request from the state table application 202. The results of the speech recognition include: the recognised recognized words of the utterance; a recognition score for each word; an actual duration for each word; and an ideal duration for each word as used in the speech recognition model. Utterance database 208 receives the results of speech recognition from voice custom server 206 and further processing is performed on the results by duration custom server 204. An example of the results and further processing is shown in Figures 5A and 5B. Duration custom server 204 acquires the data in utterance database 208 and compares the spoken duration of the actual word in an utterance with the ideal duration, this is further described with reference to duration custom server method 400 of Figure 4.

[0032] Referring to Figure 3, method 300 performed by state table application 202 is described in more detail. The first step is acquiring an utterance (step 302) from a user connected to the IVR 100 after prompting the user to speak into the telephone. Recognising Recognizing a word string from an utterance (step 304) is performed

through the custom voice server 206 using the voice server 116 and speech recognition engine 118. The results of the recognition are placed into the utterance database 208. Step 306 calculates a duration ratio. A comparison of the actual duration of utterance with an ideal duration of recognised recognized words is performed by duration custom server 204 by acquiring the values from the utterance database 208. The duration custom server also calculates an average recognition score for the whole utterance which is computed using an average of the recognition scores for all of the words. In step 308, the state table application 202 prompts the user with "please speak a little faster next time" or "please speak a little slower next time" depending on the duration ratio. In this example a duration ratio of more than one is an indication that the user is speaking slower than the ideal speed. A duration ratio of less than one indicates that the user is speaking faster than the ideal speed. The application then re-acquires the utterance (step 310) if there are words with recognition scores below a lower threshold recognition score, that is below 60%. A lower threshold recognition score is different for each speech recognition engine and configuration of the engine so, by way of example only, 60% is taken as the lower threshold recognition score to explain the embodiment. If there are words in the utterance database with recognition scores below 60% then the application re-acquires the utterance at step 314. Otherwise the method finishes at step 312 and continues with the remainder of the state table application accepting or rejecting the recognised recognized words. In normal operation re-acquisition is only performed once or twice and the best result is used or the result is negated. Step 316 skips prompt step 308 and the re-acquire step 310 if there is no need to prompt the user to speak slower or quicker. This situation occurs when the duration ratio is within a de minimus value, for example between 1.2 and 0.80 but also when the overall recognition value is above an upper threshold recognition score, for example 90%.

[0033] Referring to Figure 4, method 400 of the duration custom server 204 is described. The state table application 202 calls the duration custom server 204 after an utterance has been recorded to compare actual duration of the utterance with an ideal duration of the recognised words (step 402). The actual duration in seconds is acquired for the first word (step 404) from the utterance database 208. Then the ideal duration in seconds for the first word is acquired (step 406) from the utterance database 208. The recognition score for the first word is acquired (step 408) from the utterance database 208. If the individual recognition score for the word is greater than the lower threshold recognition score (60%) then the duration ratio is calculated (step 410) by dividing the actual duration by the ideal duration. If the word is not the last word then the process re-starts at step 404 with the next word in the utterance (step 412). If the word is the last word then an average duration ratio is calculated for words with a recognition score above the lower threshold recognition score (step 414). Method 400 ends at step 416.

[0034] The tables in Figure 5A and Figure 5B are example utterance sets of words as stored in the utterance database according to a preferred embodiment of the present invention. Referring to Figure 5A there is shown table 500 including: recognised words in column A; a recognition score for each word in column B; the actual duration of each word as estimated by the recognition engine 118 in column C; the ideal duration of each word is modeled by the recognition engine 118 in column D; and the duration ratio as calculated by the duration custom server, in column E. Cell B6 of table 500 is the average recognition score calculated by taking an average of all the individual recognition scores. Cell E6 of table 500 is the average of all the duration ratios with acceptable recognition scores as calculated by the duration custom server.

[0039] The user's utterance is sent to the speech recognition engine which creates recognition scores and durations for each ~~recognised~~ recognized word (table 500). One word at a time the recognition scores (table 500 column B) are checked, and all found to be above the lower threshold recognition score 60%. Their actual durations (table 500 column C) compared to that of ideal word durations (table 500 column D) to produce individual duration ratio (table 500 column E). The average of the duration ratios is shown in table 500 Cell B6 and is less than one at 0.78. This means that the actual utterance was shorter than the ideal and the user is speaking quicker than the ideal. All the words were successfully ~~recognised~~ recognized (recognition score above 60%) and the average duration ratio is less than the de minimus value of 0.80 so the application prompts the user to speak more slowly next time. Since the all the words were successfully ~~recognised~~ recognized the application does not re-acquire the utterance at this time.

[0041] The user's utterance is sent to the speech recognition engine which creates recognition scores and durations for each ~~recognised~~ recognized word and places them in the utterance database (table 502). One word at a time the recognition score is checked and the score for the first word is found to be lower than the lower threshold recognition score. Ignoring this first word, the duration ratio for the actual and ideal durations for the remaining three words are found (table 502 column E). This time the actual duration for saying the three words is greater than the ideal duration and this indicates that the user is speaking too slowly.

[0042] From table 502 the duration ratios of the actual duration (Column C) and the ideal duration (column D) for the last three words (column C) is 2.1, 1.9, and 1.7 (column E) which averages 1.9 (Cell E6). Therefore the actual duration for the

reeognised recognized words is greater than the ideal duration and the user is speaking slower than the ideal. Since the recognition score for the first word[[s]] is below the lower threshold recognition score (60%) then re-acquisition of the utterance is necessary.

[0044] The user's re-acquired utterance is sent to the speech recognition engine which creates recognition scores and durations for each reeognised recognized word as before. If the recognition score for the individual words are above the lower threshold recognition score (60%) then the application continues as normal with the rest of the voice application.

[0045] Although the embodiment has been described in terms of an IBM™ IVR for AIX™ other IVR's can be used to implement the invention. For instance IBM WebSphere™ Voice Response for Windows[[*]] NT[[*]]™ and Windows 2000™ with DirectTalk™ Technology is an interactive voice response (IVR) product that is for users who prefer a Windows™-based operating environment to run self-service applications. WebSphere™ [[v]]Voice Response is capable of supporting simple to complex applications and can scale to thousands of lines in a networked configuration. [[*]]Windows™, Windows 2000™ and Windows NT™ are trademarks of Microsoft Corporation in the United States, other countries, or both.

[0049] In summary there is disclosed an interactive voice response system, method and computer program product for prompting a user with speech speed feedback during speech recognition. A user who speaks too slowly or too quickly may speak even more slowly or quickly in response to an error in speech recognition. The present system aims to give the user feedback on the speed of speaking. The method includes: acquiring an utterance from a user; reeognising recognizing a string of words from the utterance;

acquiring for each word the ratio of actual duration of delivery to ideal duration; calculating an average ratio for all the words wherein the average ratio is an indication of the speed of the delivery of the utterance; and prompting the user as to the speed of delivery of the utterance according to the average ratio.